UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

LEVEL II

1. REPORT NUMBER: AFOSR-TR- 80-0266

2. GOVT ACCESSION NO.

3. RECIPIENT'S CATALOG NUMBER

4. TITLE (and Subtitle): RESEARCH ON THE EVALUATION OF COMPUTER SYSTEM ARCHITECTURES

5. TYPE OF REPORT & PERIOD COVERED: Interim rept.

6. PERFORMING ORG. REPORT NUMBER

7. AUTHOR(s): Ashok K. Agrawala

8. CONTRACT OR GRANT NUMBER(s): AFOSR-78-3654

9. PERFORMING ORGANIZATION NAME AND ADDRESS:
University of Maryland
Department of Computer Science
College Park, MD 20742

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 61102F 2304/A2

11. CONTROLLING OFFICE NAME AND ADDRESS:
Air Force Office of Scientific Research/NM
Bolling AFB, Washington, DC 20332

12. REPORT DATE: Feb[illegible]80

13. NUMBER OF PAGES: Eight

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)

15. SECURITY CLASS. (of this report): UNCLASSIFIED

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

80 4 7 150

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This study of Computer System Architecture Evaluation is divided in four subproblems. Any evaluation of a system architecture has to be done with reference to the processing demands placed on it. The first subproblem we have been studying is that of suitable representation of the user processing needs. Given a suitable characterization of the user processing needs, various techniques may be used for the system architecture evaluation. The second part of this study is aimed at the techniques applicable for architecture evaluation. The usefulness of the approaches to the first two subproblems can be assessed by

20. Abstract cont.

real system examples. Several such examples have been undertaken in ths project.

The mathematical framework available for the study of system architectures is inadequate. We are formulating a hierarchical framework to study the system architectures at various levels.

The rest of this report presents the accomplishments to date in the solutions of the four subproblems, namely:

1. Problem Domain Characterization
2. Evaluation Techniques
3. Case Studies
4. Framework for study of architectures

Most of our results have been published as the technical reports of the Department of Computer Science and some have been submitted/accepted for publication in Journals and Proceedings of Conferences. Copies of the published papers and reports are enclosed with this report.

AFOSR-TR- 80-0266

FEBRUARY 1980

RESEARCH ON THE EVALUATION
OF
COMPUTER SYSTEM ARCHITECTURES

Interim Scientific Report on the work supported by the grant No. AFOSR-78-3654 from the Air Force Office of Scientific Research to the Department of Computer Science, University of Maryland, College Park, MD.

Principal Investigator: Dr. Ashok K. Agrawala.

80 4 7 150

Introduction

The problem of design and evaluation of system architecture has received a rather ad hoc treatment. Over the years many computer systems have been designed and implemented. Most of these designs were carried out on a heuristic basis, usually trying to create the fastest or most general system with a particular technology. The system architectures were frequently decided on the basis of the available technology rather than the computational characteristics of the processing needs of the user community. Our research efforts are aimed at developing methodologies and techniques for systems architecture evaluation.

This study of Computer System Architecture Evaluation is divided in four subproblems. Any evaluation of a system architecture has to be done with reference to the processing demands placed on it. The first subproblem we have been studying is that of suitable representation of the user processing needs. Given a suitable characterization of the user processing needs, various techniques may be used for the system architecture evaluation. The second part of this study is aimed at the techniques applicable for architecture evaluation. The usefulness of the approaches to the first two subproblems can be assessed by real system examples. Several such examples have been undertaken in this project.

The mathematical framework available for the study of system architectures is inadequate. We are formulating a hierarchical framework to study the system architectures at various levels.

The rest of this report presents the accomplishments to date in the solutions of the four subproblems, namely:

1. Problem Domain Characterization
2. Evaluation Techniques

3. Case Studies

4. Framework for study of architectures

Most of our results have been published as the technical reports of the Department of Computer Science and some have been submitted/accepted for publication in Journals and Proceedings of conferences. Copies of the published papers and reports are enclosed with this report.

Problem Domain Characterization

In evaluating the performance of computer system architectures we have to explicitly consider the workload with respect to which such an evaluation is to be done. A system suitable for numeric computations may not be well suited for business applications. In some of our studies we considered image processing as the problem domain.

In order to understand some of the characteristics of the workloads presented by the image processing activities, we examined the workload of the Computer Vision Lab on the Univac machines during November 1978. A period of 17 days was taken (November 1-November 17) and all the log tapes for this period were used in the study. A clustering based characterization of the workload was carried out. There were a total of 19 clusters of jobs. An internal memo describing the results is enclosed. The workload created by the Computer Vision Lab activities was spread in 5-6 clusters. These clusters also contained many other research jobs. The results of this experiment imply that from the resource usage stand point the image processing jobs handled in the research environment of a University do not look that different from many other research jobs.

The activities of the Computer Vision Lab (CVL) not only include actual processing of images but also developing techniques for image processing. Therefore, a significant share of the workload reflects program development activity which is rather similar to program development in other fields. For many image processing tasks a dedicated Grinell system coupled to the PDP 11/45 system is used. Therefore, the CVL workload handled by the Univac machines does not reflect that activity. These causes may contribute to the findings of our experiment regarding the cluster-based characterization of the workload. At present we are examining alternate ways of characterizing problem domains which may be more suitable for architecture evaluation.

Evaluation Techniques

Network of queues models have been sucessfully used to evaluate computer system performance. Analytic, simulation and approximation techniques exist to solve these models. We examined the applicability of queueing network models to system architecture evaluation problem.

The network of queues models for computer systems make strong assumptions regarding the way various resources are used by jobs and tasks, and the kind of interactions that exist among the jobs, tasks and resources. While many new generalizations have been reported in the literature, the models which can be solved today (using exact or approximation techniques) do not adequately reflect the realistic constraints of the system at an architectural level. We are examining techniques for creating and solving models which reflect the computer system realities more accurately.

In problem domain characterization is is possible to come up with approximate values of workload parameters. (In any parameter estimation problem this is a common issue.) We have investigated the sensitivity of the parameters of queueing network models on the performance of the system. Theoretical bounds relating the errors in parameter estimation to the errors in performance quantities were obtained.

Nearly all the developments in the queueing theory have been aimed at the steady state behavior of queues. In practice, however, transient response is frequently necessary. Recently we developed techniques which allow the computation of the transient response characteristics of a single server queue. A paper on this technique has been accepted for publication in the International Journal of Information Sciences.

The approach taken in obtaining a transient solution of a single server queue is to evaluate the virtual waiting time for the server. Based on this approach more complex queueing systems can also be studied. Recently we obtained a solution to a queueing system with exponential queue with cyclic arrival patterns.

While these developments are theoretical in nature they have direct and immediate applications. Based on the analysis of the cyclic queues we determined that corresponding to a random routing there exists a deterministic routing which is superior. An adaptive routing technique which takes into account the transient behavior of the queues was also formulated.

The routing decisions which minimize the average time in the system are also effected by future arrivals. Only probabilistic information regarding their arrival may be available at their arrival instant. An approach to routing which

In problem domain characterization is is possible to come up with approximate values of workload parameters. (In any parameter estimation problem this is a common issue.) We have investigated the sensitivity of the parameters of queueing network models on the performance of the system. Theoretical bounds relating the errors in parameter estimation to the errors in performance quantities were obtained.

Nearly all the developments in the queueing theory have been aimed at the steady state behavior of queues. In practice, however, transient response is frequently necessary. Recently we developed techniques which allow the computation of the transient response characteristics of a single server queue. A paper on this technique has been accepted for publication in the International Journal of Information Sciences.

The approach taken in obtaining a transient solution of a single server queue is to evaluate the virtual waiting time for the server. Based on this approach more complex queueing systems can also be studied. Recently we obtained a solution to a queueing system with exponential queue with cyclic arrival patterns.

While these developments are theoretical in nature they have direct and immediate applications. Based on the analysis of the cyclic queues we determined that corresponding to a random routing there exists a deterministic routing which is superior. An adaptive routing technique which takes into account the transient behavior of the queues was also formulated.

The routing decisions which minimize the average time in the system are also effected by future arrivals. Only probabilistic information regarding their arrival may be available at their arrival instant. An approach to routing which

Framework for Studies of Computer System Architectures

In using a computer system to solve a problem, the original user problem has to undergo several mappings or translations. At each step of this process a level of detail of the computer system gets involved. Today no framework exists which allows us to adequately model the system and its workload at various levels. We are attempting to formulate a hierarchical framework which will allow the system architecture to be studied at various levels of detail.

The approach taken in this work is to consider the computer system architecture in terms of a triple consisting of (algortihm, structure and implementation). The algorithm is the term used here to capture the problem domain characteristics relevant to the system architecture. The structure specifies the basic structure of the system. The implementation is the mapping of the structure to the physical system. The work on this framework started recently and a technical report documenting some of our initial findings is in preparation.

Concluding Remarks

Our studies in the evaluation of computer system architectures have attempted to examine the major subproblems involved. In the earlier sections the major accomplishments have been summarized. We have published these results as technical reports. Several of these are also getting published in journals and conference proceedings. A copy of the technical reports is enclosed here.

Enclosures

1. "Computer Performance Predication Via Analytical Modeling - An Experiment", Proc. of 1979 Conference on Sim., Meas, and Modeling of Computer Systems, August 1979, (Dowdy, Agrawala, Gordon, Tripathi).,

2. "Transient Solution of the Virtual Waiting Time of a Single Server Queue and its Applications", TR-806, August 1979 - To appear in Information Sciences, April 1980, (Agrawala, Tripathi).

3. "Some Strategies for the Multiple Producer/Multiple Consumer Problem", TR-816, October 1979, (Ricart, Agrawala).

4. "Adaptive Routing Using a Virtual Waiting Time Technique", TR-817, November 1979, (Agrawala, Tripathi, Ricart).

5. "Analysis of Design Alternatives for a Packet Switched I/O System", TR-830, November 1979 -- Accepted for Performance '80, (Upton, Tripathi).

6. "On Multiprogramming Level in Queueing Networks", TR-831, November 1979, (Dowdy, Gordon, Agre).

7. "The Impact of Certain Parameter Estimation Errors in Queueing Network Models", TR-832, November 1979 -- Accepted for Performance '80, (Gordon, Dowdy).

8. "Formal Approaches to Protocol Design and Implimentations: A Survey", TR-840, December 1979, (Thareja, Agrawala, Larsen).

9. "Characterizing the Workload of the University of Maryland Computer Science Center", Internal Memo 7912-01, December 1979, (Agrawala, Gordon, Mohr).

10. "On an Exponential Server with General Cyclic Arrivals", TR-867, February 1980, (Agrawala, Tripathi).

11. "The Effect of the Future in Work Distribution", TR-868, February 1980, (Ricart, Agrawala).